Building a Balanced k-d Tree in O(kn log n) Time

Russell A. Brown

arXiv:1410.5420v27 [cs.DS] 7 May 2015

ABSTRACT

The original description of the k-d tree recognized that rebalancing techniques, such as are used to build an AVL tree or a red-black tree, are not applicable to a k-d tree. Hence, in order to build a balanced k-d tree, it is necessary to find the median of the data for each recursive subdivision of those data. The sort or selection that is used to find the median for each subdivision strongly influences the computational complexity of building a k-d tree.

This paper discusses an alternative algorithm that builds a balanced k-d tree by presorting the data in each of k dimensions prior to building the tree. It then preserves the order of these k sorts during tree construction and thereby avoids the requirement for any further sorting. Moreover, this algorithm is amenable to parallel execution via multiple threads. Compared to an algorithm that finds the median for each recursive subdivision, this presorting algorithm has equivalent performance for four dimensions and better performance for three or fewer dimensions.

1. INTRODUCTION

Bentley introduced the k-d tree as a binary tree that stores k-dimensional data [4]. Like a standard binary tree, the k-d tree subdivides data at each recursive level of the tree. Unlike a standard binary tree that uses only one key for all levels of the tree, the k-d tree uses k keys and cycles through these keys for successive levels of the tree. For example, to build a k-d tree from three-dimensional points that comprise (x, y, z) coordinates, the keys would be cycled as $x, y, z, x, y, z...$ for successive levels of the k-d tree. A more elaborate scheme for cycling the keys chooses the coordinate that has the widest dispersion or largest variance to be the key for a particular level of recursion [8].

Due to the use of different keys at successive levels of the tree, it is not possible to employ rebalancing techniques, such as are used to build an AVL tree [1] or a red-black tree [3, 10], when building a k-d tree. Hence, the typical approach to building a balanced k-d tree finds the median of the data for each recursive subdivision of those data. Bentley showed that if the median of n elements were found in $O(n)$ time, it would be possible to build a depth-balanced k-d tree in $O(n \log n)$ time. However, algorithms that find the median in guaranteed $O(n)$ time are somewhat complicated [6, 7]. Quicksort [13] finds the median in $O(n)$ time in the best case, but in $O(n^2)$ time in the worst case [17]. Merge sort [9] and heap sort [16] find the median in guaranteed $O(n \log n)$ time, which leads to $O(n \log^2 n)$ time for building a balanced k-d tree [15].

An alternative approach to building a balanced k-d tree presorts the data prior to building the tree and avoids resorting for each recursive subdivision. Two such algorithms have been reported that sort triangles for three-dimensional graphics ray tracing and that have best-case complexity of $O(n \log n)$ but undetermined worst-case complexity [12, 15]. The algorithm that is described in the present article presorts points in each of k dimensions prior to building the k-d tree, then maintains the order of these k sorts when building a balanced k-d tree and thereby achieves a worst-case complexity of $O(kn \log n)$.

2. IMPLEMENTATION

2.1 The $O(kn \log n)$ Algorithm

Consider the 15 (x, y, z) tuples that are stored in elements 0 through 14 of the "Tuples" array that is shown at the left side of Figure 1. The k-d tree-building algorithm begins by presorting the tuples in their x-, y- and z-coordinates via three executions of merge sort. These three sorts do not in fact sort the x-, y- and z-coordinates by using these coordinates as sort keys entirely independently of one another; instead, x, y and z form the most significant portions of the respective super keys x:y:z, y:z:x and z:x:y that represent cyclic permutations of x, y and z. The symbols for these super keys use a colon to designate the concatenation of the individual x, y and z values. Hence, for example, the symbol z:x:y represents a super key wherein z is the most significant portion of the super key, x is the middle portion of the super key and y is the least significant portion of the super key.

The merge sorts employ super keys, instead of keys that are merely the individual x-, y- and z-coordinates, in order to detect and remove duplicate tuples, as will be explained later. The merge sorts do not reorder the tuples array; rather, they reorder three index arrays whose elements are indices into the tuples array. The initial order of these index arrays is established by the merge sorts and is shown in Figure 1 in the xyz, yzx and zxy columns under "Initial Indices". In this figure, xyz, yzx and zxy are shorthand notations for the super keys x:y:z, y:z:x and z:x:y, respectively.

The xyz, yzx and zxy columns under "Initial Indices" represent the initial order of the xyz-, yzx- and zxy-index arrays that indicate the results of the three merge sorts. For example, elements 0, 1, ... 13, 14 of the xyz-index array

contain the sequence 11, 13, ... 2, 8 that represents the respective tuples $(1, 6, 8)$; $(2, 1, 3)$; ... $(9, 6, 7)$; $(9, 7, 8)$ that were ordered via merge sort using the respective super keys 1:6:8, 2:1:3, ... 9:6:7, 9:7:8. Similarly, elements 0, 1, ... 13, 14 of the yzx-index array contain the sequence 13, 4, ... 8, 3 that represents the respective tuples $(2, 1, 3)$; $(8, 1, 5)$; ... $(9, 7, 8)$; $(4, 7, 9)$ that were ordered via merge sort using the respective super keys 1:3:2, 1:5:8, ... 7:8:9, 7:9:4. Lastly, elements 0, 1, ... 13, 14 of the zxy-index array contain the sequence 9, 6, ... 8, 3 that represents the respective tuples $(6, 3, 1)$; $(9, 4, 1)$; ... $(9, 7, 8)$; $(4, 7, 9)$ that were ordered via merge sort using the respective super keys 1:6:3, 1:9:4, ... 8:9:7, 9:4:7.

	Tuples (x,y,z)	Initial Indices xyz	yzx	zxy	After First Split xyz	yzx	zxy	After Second Split xyz	yzx	zxy
0	(2,3,3)	11	13	9	11	13	9	13	13	9
1	(5,4,2)	13	4	6	13	9	1	0	9	13
2	(9,6,7)	0	5	1	0	0	13	9	0	0
3	(4,7,9)	10	9	7	10	1	0			
4	(8,1,5)	3	0	13	3	10	10	11	10	10
5	(7,2,6)	1	6	0	1	11	11	10	11	11
6	(9,4,1)	9	1	12	9	3	3	3	3	3
7	(8,4,2)	5	7	10						
8	(9,7,8)	4	10	4	4	4	6	4	4	6
9	(6,3,1)	7	12	5	7	6	7	7	6	7
10	(3,4,5)	14	2	14	14	7	12	6	7	4
11	(1,6,8)	6	11	2	6	12	4			
12	(9,5,3)	12	14	11	12	2	14	14	2	14
13	(2,1,3)	2	8	8	2	14	2	2	14	2
14	(8,7,6)	8	3	3	8	8	8	8	8	8

Figure 1: An (x, y, z) tuple array and xyz-, yzx- and zxy-index arrays.

The next step of the k-d tree-building algorithm partitions the (x, y, z) tuples in x using the x:y:z super key that is specified by the median element of the xyz-index array under "Initial Indices". This median element is located at address 7 of this array; its value is 5 and specifies the tuple $(7, 2, 6)$ for which the x:y:z super key is 7:2:6. The partitioning does not reorder the tuples array; instead, it reorders the yzx- and zxy-index arrays. The xyz-index array requires no partitioning because it is already sorted in x. However, the yzx- and zxy-index arrays require partitioning in x using the x:y:z super key 7:2:6 that is specified by the median element of the xyz-index array.

This partitioning is accomplished for the yzx-index array as follows. The elements of the yzx-index array are retrieved in order of increasing address from 0 to 14. The x:y:z super key that is specified by each element of the yzx-index array is compared to the 7:2:6 super key that is specified by the median element of the xyz-index array. Each element of the yzx-index array is copied to either the lower or upper half of a temporary index array, depending on the result of this comparison. After all of the elements of the yzx-index array have been processed in this manner, the temporary index array replaces the yzx-index array and becomes the new yzx-index array that is depicted in Figure 1 under "After First Split." The partitioning of the first six elements of the yzx-index array is discussed below and provides insight into the details of the k-d tree-building algorithm.

The element at address 0 of the yzx-index array is 13 and specifies the tuple $(2, 1, 3)$ for which the x:y:z super key is 2:1:3. This super key is less than the median super key 7:2:6; hence, the element at address 0 of the yzx-index array is copied to address 0 in the new yzx-index array, which is the lowest address in the lower half of the new yzx-index array. The element at address 1 of the yzx-index array is 4 and specifies the tuple $(8, 1, 5)$ for which the x:y:z super key is 8:1:5. This super key is greater than the median super key 7:2:6; hence, the element at address 1 of the yzx-index array is copied to address 8 in the upper half of the new yzx-index array, which is the lowest address in the upper half of the new yzx-index array. The element at address 2 of the yzx-index array is 5 and specifies the tuple $(7, 2, 6)$ for which the x:y:z super key is 7:2:6. This super key equals the median super key 7:2:6; hence, the element at address 2 in the yzx-index array is ignored and not copied to the new yzx-index array.

The element at address 3 of the yzx-index array is 9 and specifies the tuple $(6, 3, 1)$ for which the x:y:z super key is 6:3:1. This super key is less than the median super key 7:2:6; hence, the element at address 3 of the yzx-index array is copied to address 1 in the lower half of the new yzx-index array, which is the second lowest address in the lower half of the new yzx-index array. The element at address 4 of the yzx-index array is 0 and specifies the tuple $(2, 3, 3)$ for which the x:y:z super key is 2:3:3. This super key is less than the median super key 7:2:6; hence, the element at address 4 of the yzx-index array is copied to address 2 in the lower half of the new yzx-index array, which is the third lowest address in the lower half of the new yzx-index array. The element at address 5 of the yzx-index array is 6 and specifies the tuple $(9, 4, 1)$ for which the x:y:z super key is 9:4:1. This super key is greater than the median super key 7:2:6; hence, the element at address 5 of the yzx-index array is copied to address 9 in the upper half of the new yzx-index array, which is the second lowest address in the upper half of the new yzx-index array.

Partitioning continues for the remaining eight elements of the yzx-index array in the manner that is described above. The partitioning of the first six elements of the yzx-index array reveals that the yzx-index array has been partitioned in x relative to the median element of the xyz-index array; this partitioning preserves the initial merge-sorted order in y within the lower and upper halves of the new yzx-index array.

Next, the zxy-index array is partitioned in x relative to the median element of the xyz-index array, which preserves the initial merge-sorted order in z for the lower and upper halves of the new zxy-index array. The reader is encouraged to audit the partitioning of the first few elements of the zxy-index array under "Initial Indices" in order to verify that these elements are correctly assigned to the lower and upper halves of the new zxy-index array that is shown in

Figure 1 under "After First Split." Because the partitioning in x preserves the initial merge-sorted order for the lower and upper halves of the yzx- and zxy-index arrays, there is no requirement for any further sorting after the k initial merge sorts.

Inspection of the lower and upper halves of the new xyz-, yzx- and zxy-index arrays in Figure 1 under "After First Split" reveals that the index value 5 is absent from the lower and upper halves of these index arrays. This value is absent from these index arrays because it is the value of the median element of the xyz-index array that specified the x:y:z super key 7:2:6 relative to which the yzx- and zxy-index arrays were partitioned in x. In order to record this partitioning, a reference to the tuple $(7, 2, 6)$ is stored in the root of the nascent k-d tree, as shown in Figure 2.

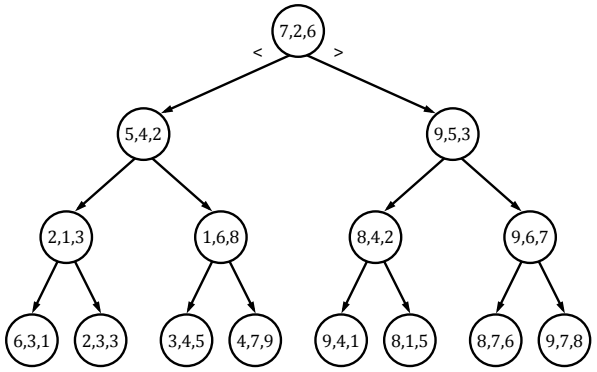

Figure 2: A k-d tree that is built from the (x, y, z) tuples of Figure 1.

Next, the lower and upper halves of the xyz-, yzx- and zxy-index arrays are processed recursively and partitioned in y to create the "less than" and "greater than" subtrees of the root of the k-d tree. Consider the lower half of the yzx-index array that is depicted in Figure 1 under "After First Split." The median element of this array is located at address 3; its value is 1 and specifies the tuple $(5, 4, 2)$ for which the y:z:x super key is 4:2:5. The lower half of the yzx-index array is already sorted in y. However, the lower halves of the zxy- and xyz-index arrays require partitioning in y relative to the y:z:x super key 4:2:5 that is specified by the median element of the lower half of the yzx-index array. The reader is encouraged to verify the result of this partitioning by inspection of the first and second fourths of the new xyz-, yzx- and zxy-index arrays that are depicted in Figure 1 under "After Second Split." The upper halves of the zxy- and xyz-index arrays are partitioned in a similar manner relative to the y:z:x super key 5:3:9 that is formed from the tuple $(9, 5, 3)$ that is specified by the value 12 of the median element at address 11 of the upper half of the yzx-index array. References to the tuples $(5, 4, 2)$ and $(9, 5, 3)$ are stored in the "less than" and "greater than" children of the root of the nascent k-d tree, as shown in Figure 2.

Recursion terminates when an index array comprises one, two or three elements. In the case of one element, a reference to the corresponding tuple is stored in a new node of the k-d tree. For two or three elements, the elements are already in sorted order in the index array, so the determination of which tuple to reference from a new node of the k-d tree and which tuple or tuples to reference from children of that node is trivial. For example, consider the four fourths of the zxy-index arrays under "After Second Split" in Figure 1. Each fourth comprises three elements, so recursion terminates. The tuples $(2, 1, 3)$ and $(1, 6, 8)$ that are specified respectively by the median elements 13 and 11 at addresses 1 and 5 of the zxy-index array are referenced by the respective "less than" and "greater than" children of node $(5, 4, 2)$ of the nascent k-d tree. Similarly, the tuples $(8, 4, 2)$ and $(9, 6, 7)$ that are specified respectively by the median elements 7 and 2 at addresses 9 and 13 of the zxy-index array are referenced by the respective "less than" and "greater than" children of node $(9, 5, 3)$ of the nascent k-d tree. The children and grandchildren of nodes $(5, 4, 2)$ and $(9, 5, 3)$ are shown in Figure 2.

The foregoing discussion reveals that the k-d tree includes "less than" and "greater than" children but no "equal" children. For this reason, duplicate (x, y, z) tuples must be removed from the data prior to building the k-d tree. After the k initial merge sorts have reordered the xyz-, yzx- and zxy-index arrays, each index array is traversed once in order to discard all but one index from each set of contiguous indices that reference identical (x, y, z) tuples. In order that adjacent indices reference identical (x, y, z) tuples, the k initial merge sorts employ x:y:z, y:z:x and z:x:y super keys instead of keys that are merely the individual x-, y- and z-coordinates. If the k initial merge sorts employed keys that were only the individual x-, y- and z-coordinates, adjacent indices within an index array could reference non-identical (x, y, z) tuples for which one or more, but not all, of the x-, y- and z-coordinates were identical. The x:y:z, y:z:x and z:x:y super keys guarantee that each group of identical (x, y, z) tuples is indexed by a set of contiguous indices within each index array. These super keys enable the removal of duplicate (x, y, z) tuples via one pass through each index array that discards adjacent indices that reference identical (x, y, z) tuples.

It is possible to optimize the use of the temporary index array such that only one temporary index array is required and such that the xyz-, yzx- and zxy-index arrays may be reused to avoid allocation of new index arrays at each level of recursion. This optimization operates as follows. The xyz-index array is copied to the temporary index array. Then the yzx-index array is partitioned in x and the result is stored in the two halves of the xyz-index array. Next, the zxy-index array is partitioned in x and the result is stored in the two halves of the yzx-index array. Finally, the temporary index array is copied to the zxy-index array. This optimization permutes the xyz-, yzx- and zxy-index arrays cyclically at each level of recursion, as is required to cycle the keys in the order $x, y, z, x, y, z...$ for successive levels of the k-d tree. Moreover, it guarantees that the x:y:z, y:z:x or z:x:y super key that is required for partitioning at a particular level of the k-d tree is always specified by the median element of the xyz-index array. The computational cost of this index array optimization is the copying of one additional index array at each level of recursion.

Recursive partitioning occurs for $\log_2(n)$ levels of the nascent k-d tree. The computational complexity of this k-d tree-building algorithm includes a $O(kn \log n)$ term for the k initial merge sorts plus a $O((k+1)\, n \log n)$ term for copying n elements of $k+1$ index arrays at each of the $\log_2(n)$ levels of recursion. This $O(kn \log n)$ k-d tree-building algo-

rithm requires storage for a tuples array of n k-dimensional tuples, plus an n-element temporary array, plus k n-element index arrays. The tuples array is immutable. The index and temporary arrays are ephemeral and are no longer required after construction of the k-d tree.

2.2 Parallel Execution

The merge-sorting function [14] and the $O(kn \log n)$ k-d tree-building function both subdivide index arrays and process non-overlapping halves of each index array via recursive calls to these functions. Hence, these functions (or Java methods) are amenable to parallel execution via multiple threads that occurs as follows.

One thread executes a recursive call of the method; this thread is designated as the parent thread. The parent thread subdivides one or more index arrays, then calls the method recursively to process the lower and upper halves of each index array. The parent thread does not execute the recursive call that processes the lower half of each index array; instead, it launches a child thread to execute that recursive call. After launching the child thread, the parent thread executes the recursive call that processes the upper half of each index array, then waits for the child thread to finish execution of the recursive call that processes the lower half of each index array.

For a balanced k-d tree, the number of threads q that are required by this thread-allocation strategy is $q = 2^d \approx n/2$ where d represents the deepest level of recursion and n represents the number of tuples. A large number of tuples would require a prohibitively large number of threads; hence, child threads are launched to only the maximum level of recursion that is allowed by the number of available threads. Beyond this maximum level of recursion, the parent thread processes both halves of each index array.

Two threads permit launching a child thread at the first level of recursion. Four threads permit launching child threads at the first two levels of recursion. Eight threads permit launching child threads at the first three levels of recursion, *et cetera*. Because threads are launched at the lowest levels of recursion, each thread processes the maximum possible workload. Because the index arrays are subdivided by their median elements at each level of recursion, all threads share the workload equally.

A disadvantage of this thread allocation strategy is that it limits the number of threads to an integer power of two. Because the level of recursion determines the number of threads, it is not possible to employ, for example, three or ten threads. An advantage of this thread allocation strategy is that it is simple and robust because synchronization involves only a parent thread and one child thread.

2.3 Results for the $O(kn \log n)$ Algorithm

The $O(kn \log n)$ k-d tree-building algorithm was implemented in the Java language, and the single-threaded performance of the merge sorting, duplicate tuple removal and k-d tree-building methods was measured using a 2.3 GHz Intel i7 processor. Figure 3 shows the total time in seconds that was required to perform the initial merge sorts, remove the duplicate tuples and build the k-d tree, plotted versus $n \log_2(n)$ for $2^{18} \leq n \leq 2^{24}$ (x, y, z, w) tuples of randomly-generated 32-bit integers. The dashed line of Figure 3 shows the least-squares fit of the total time t to the function $t = mn \log_2(n)$ where m is the slope of the line. The correlation coefficient $r = 0.998$ indicates an adequate least-squares fit; hence, the execution times are proportional to $n \log_2(n)$. The question of whether these execution times are proportional to k will be addressed in Section 4 of this article.

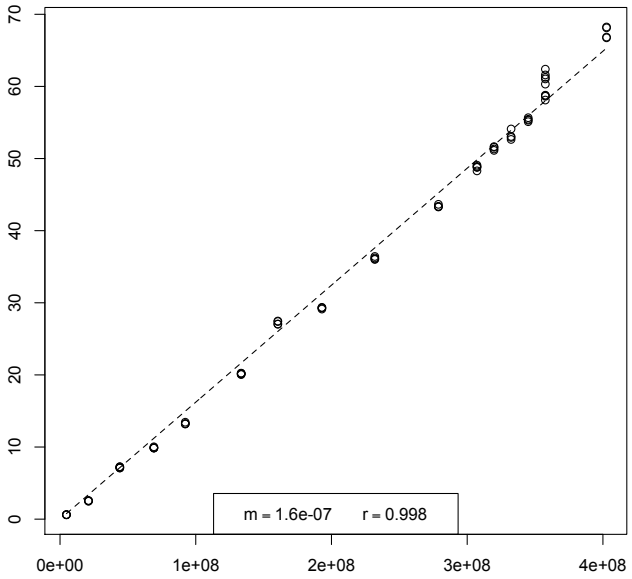

Figure 3: The total of merge sorting, duplicate tuple removal and k-d tree-building times (seconds) is plotted vs. $n \log_2(n)$ for the application of the $O(kn \log n)$ k-d tree-building algorithm to $2^{18} \leq n \leq 2^{24}$ (x, y, z, w) tuples of randomly-generated 32-bit integers.

The $O(kn \log n)$ k-d tree-building algorithm was parallelized via Java threads and its performance was measured for one to eight threads using a 2.3 GHz Intel quad-core i7 processor. Figure 4 shows the total time in seconds that was required to perform the initial merge sorts, remove the duplicate tuples and build the k-d tree, plotted versus the number of threads q for $n = 2^{24}$ (x, y, z, w) tuples of randomly-generated 32-bit integers. The dashed curve of Figure 4 shows the least-squares fit of the total time t to the equation

$$t = t_s + \frac{t_1}{q} + m_c(q - 1) \tag{1}$$

This equation will be discussed in Section 4 of this article. The correlation coefficient $r = 0.9867$ indicates an acceptable least-squares fit.

3. COMPARATIVE PERFORMANCE

3.1 The $O(n \log n)$ Algorithm

In order to understand the performance of the $O(kn \log n)$ k-d tree-building algorithm relative to that of other algorithms, it was compared to a $O(n \log n)$ k-d tree-building algorithm that incorporates a $O(n)$ median-finding algorithm [6, 7]. Most of the complexity of the $O(n \log n)$ algorithm is limited to the $O(n)$ median-finding algorithm. The application of this $O(n \log n)$ k-d tree-building algorithm to sort (x, y, z) tuples is described as follows.

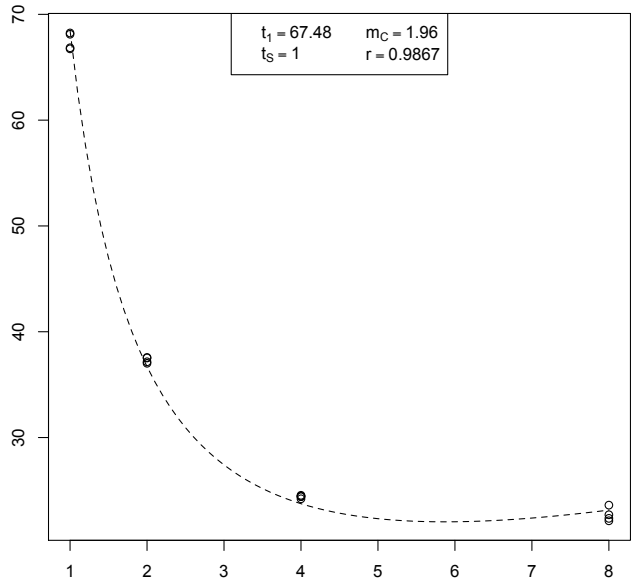

Figure 4: The total of merge sorting, duplicate tuple removal and k-d tree-building times (seconds) is plotted vs. the number of threads for the application of the $O(kn \log n)$ k-d tree-building algorithm to $n = 2^{24}$ (x, y, z, w) tuples of randomly-generated 32-bit integers.

First, an index array is created and merge sorted in x, y or z via one of the x:y:z, y:z:x and z:x:y super keys; the choice of super key is arbitrary. The initial merge sort does not reorder the tuples array; instead, it reorders the index array whose elements are indices into the tuples array. Next, duplicate (x, y, z) tuples are removed via one pass through the index array, as discussed in Section 2.1 of this article.

The subsequent k-d tree-building step partitions the index array recursively. At each level of recursion, the median element of the index array is found in $O(n)$ time using the x:y:z, y:z:x or z:x:y super key that is appropriate to that level of recursion. A convenient feature of the $O(n)$ median-finding algorithm is that the index array is partitioned relative to the median element during the search for the median element. Hence, once the median element has been found, a reference to the (x, y, z) tuple that the median element specifies is stored in the root of the nascent k-d tree, as shown in Figure 2. The lower and upper halves of the index array are processed recursively to create the "less than" and "greater than" subtrees of the root of the k-d tree. The $O(n \log n)$ k-d tree-building method processes non-overlapping halves of the index array via recursive calls to this method. Hence, this method is amenable to parallel execution via multiple threads in the manner that was explained in Section 2.2 of this article.

Recursion terminates when the index array comprises one, two or three elements. In the case of one element, a reference to the corresponding tuple is stored in a new node of the k-d tree. For two elements, a reference to the tuple that corresponds to the first element is stored in a new node of the k-d tree, then the super keys of the two elements are compared to decide whether to reference the tuple that corresponds to the second element from the "less than" or "greater than" child of that node. For three elements, the index array is sorted via insertion sort [5] to determine which tuple to reference from a new node of the k-d tree and which tuples to reference from the children of that node.

Recursive partitioning occurs for $\log_2(n)$ levels of the nascent k-d tree. The computational complexity of this k-d tree-building algorithm includes a $O(n \log n)$ term for the initial merge sort plus another $O(n \log n)$ term for partitioning n elements of the index array at each of the $\log_2(n)$ levels of recursion. This $O(n \log n)$ k-d tree-building algorithm requires storage for a tuples array of n k-dimensional tuples, plus an n-element index array, plus an $n/2$-element temporary array. The tuples array is immutable. The index and temporary arrays are ephemeral and are no longer required after construction of the k-d tree.

3.2 Results for the $O(n \log n)$ Algorithm

The $O(n \log n)$ k-d tree-building algorithm was implemented in the Java language, and the single-threaded performance of the merge sorting, duplicate tuple removal and k-d tree-building methods was measured using a 2.3 GHz Intel i7 processor. Figure 5 shows the total time in seconds that was required to perform the initial merge sort, remove the duplicate tuples and build the k-d tree, plotted versus $n \log_2(n)$

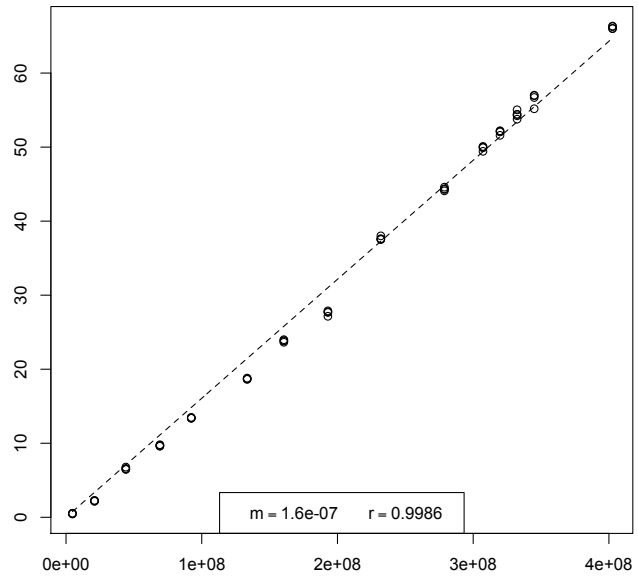

Figure 5: The total of merge sorting, duplicate tuple removal and k-d tree-building times (seconds) is plotted vs. $n \log_2(n)$ for the application of the $O(n \log n)$ k-d tree-building algorithm to $2^{18} \leq n \leq 2^{24}$ (x, y, z, w) tuples of randomly-generated 32-bit integers.

for $2^{18} \leq n \leq 2^{24}$ (x, y, z, w) tuples of randomly- generated 32-bit integers. The dashed line of Figure 5 shows the least-squares fit of the total time t to the function $t = mn \log_2(n)$ where m is the slope of the line. The correlation coefficient $r = 0.9986$ indicates an adequate least-squares fit.

The $O(n \log n)$ k-d tree-building algorithm was parallelized via Java threads and its performance was measured for one to eight threads using a 2.3 GHz Intel quad-core i7 processor. Figure 6 shows the total time in seconds that was required to

perform the initial merge sort, remove the duplicate tuples and build the k-d tree, plotted versus the number of threads q for $n = 2^{24}$ (x, y, z, w) tuples of randomly-generated 32-bit integers. The dashed curve of Figure 6 shows the least-squares fit of the total time t to Equation 1. The correlation coefficient $r = 0.9958$ indicates an acceptable least-squares fit.

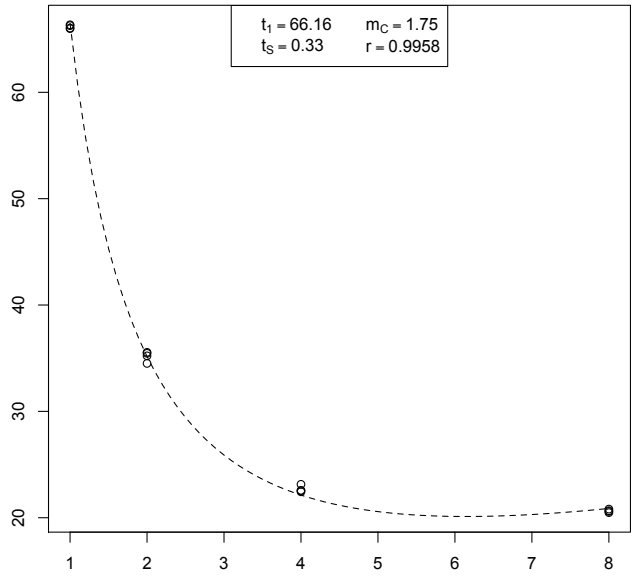

Figure 6: The total of merge sorting, duplicate tuple removal and k-d tree-building times (seconds) is plotted vs. the number of threads for the application of the $O(n \log n)$ k-d tree-building algorithm to $n = 2^{24}$ (x, y, z, w) tuples of randomly-generated 32-bit integers.

4. DISCUSSION

Figures 3 and 5 demonstrate that the execution times of the $O(kn \log n)$ and $O(n \log n)$ k-d tree-building algorithms are proportional to $n \log_2(n)$. Figures 4 and 6 show that the k-d tree-building algorithms scale for multiple threads. For either algorithm, the execution by eight threads on a Intel quad-core i7 processor, which supports concurrent execution of two threads per core, increases the execution speed by approximately three times relative to the speed of one thread. The execution time t does not adhere to the Amdahl [2] model

$$t = t_s + \frac{t_1}{q} \tag{2}$$

but rather to the model that is expressed by Equation 1 in Section 2.3 of this article

$$t = t_s + \frac{t_1}{q} + m_c (q - 1)$$

In this equation, q is the number of threads, t_s represents the time required to execute the serial or non-parallelizable portion of the algorithm, t_1 represents the time required to execute the parallelizable portion of the algorithm via one thread, and $m_c (q - 1)$ models an additional limitation to the performance of multi-threaded execution that the Amdahl model fails to capture.

This additional limitation to performance may occur due to cache misses in a shared-memory architecture. During multi-threaded execution of the k-d tree building algorithm, any thread may read from any address of the (x, y, z, w) tuples array, as directed by a specific index from an index array. This unrestricted access to the (x, y, z, w) tuples array may cause cache misses in a shared-memory architecture because two threads could simultaneously attempt to access two different tuples that map into the same cache line of a shared cache memory. The cache miss term $m_c (q - 1)$ of Equation 1 models the performance limitation, which results from cache misses, as a linear function of q [11].

Differentiating Equation 1 with respect to q yields

$$\frac{dt}{dq} = m_c - \frac{t_1}{q^2} \tag{3}$$

Setting dt/dq to zero in Equation 3 and solving for q predicts that the minimum execution time should occur at $q = \sqrt{t_1/m_c}$ threads. Substituting into this square root the respective values of t_1 and m_c that were obtained via least-squares fitting for the $O(kn \log n)$ and $O(n \log n)$ algorithms predicts minima at $q = 5.87$ and $q = 6.15$ threads, respectively. These minima are depicted in Figures 4 and 6, which predict decreased performance of both k-d tree building algorithms for greater than eight threads. The decreased performance for as few as eight threads is a consequence of the relatively large values for m_c (1.96 and 1.75 seconds per thread, respectively) that were obtained via least-squares fitting. If execution time measurements were obtained using an Intel multi-core i7 processor that supports concurrent execution of more than eight threads, those measurements should confirm decreased performance for greater than eight threads.

The similar parallelizable times t_1 for both algorithms (67.48 and 66.16 thread-seconds) indicate that their single-threaded performance is about equal. This effect is due to a fortuitous choice of $k = 4$ that specifies test data that comprise (x, y, z, w) tuples. Because the execution time of the $O(kn \log n)$ algorithm should be proportional to k but the execution time of the $O(n \log n)$ algorithm should not, these two algorithms are expected to have unequal performance for a different choice of k. In order to test this hypothesis, each algorithm was utilized to build five different k-d trees. For each k-d tree, 2^{24} k-dimensional tuples of randomly-generated 32-bit integers were created using a different value of $k = 2, 3, 4, 5, 6$. The performance of each algorithm was measured using a single thread of a 2.3 GHz Intel i7 processor. The results are shown in Figure 7.

Figure 7 shows the total time in seconds that was required to perform the initial merge sorts, remove the duplicate tuples and build the k-d tree via the $O(kn \log n)$ and $O(n \log n)$ algorithms for $n = 2^{24}$ k-dimensional tuples of randomly-generated 32-bit integers, plotted versus the number of dimensions $k = 2, 3, 4, 5, 6$. This figure demonstrates that the execution time of the $O(kn \log n)$ algorithm is proportional to k but the execution time of the $O(n \log n)$ algorithm is not. In this figure, the slope of the solid line ($m_{\text{SOLID}} = 18.07$ seconds per dimension) indicates that for 2^{24} tuples, each additional dimension increases the execution time of the $O(kn \log n)$ algorithm by 18 seconds. For $k = 4$, the two algorithms have equal performance.

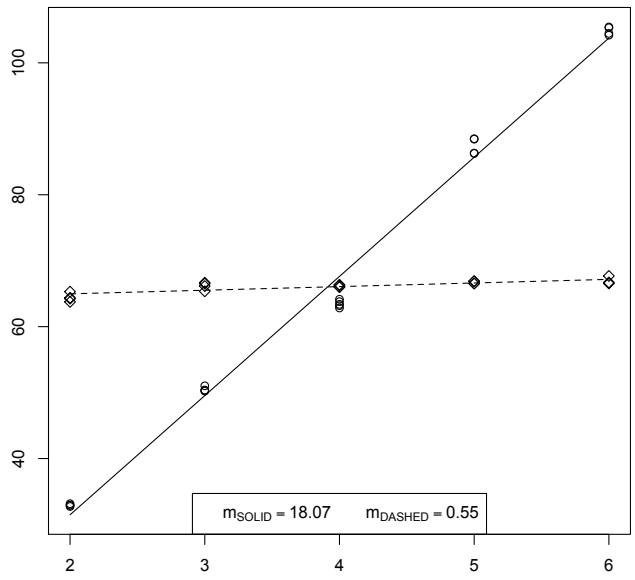

Figure 7: The total of merge sorting, duplicate tuple removal and k-d tree-building times (seconds) is plotted vs. the number of dimensions k for the application of the $O(kn \log n)$ algorithm (solid line and circles) and the $O(n \log n)$ algorithm (dashed line and diamonds) to build a k-d tree from $n = 2^{24}$ k-dimensional tuples of randomly-generated 32-bit integers.

In Figure 7, the slope of the dashed line ($m_{\text{DASHED}} = 0.55$ seconds per dimension) suggests that the execution time of the $O(n \log n)$ algorithm might be proportional to k. However, this apparent scaling is an artifact that is related to the fact that the k-dimensional tuples comprise randomly-generated 32-bit integers.

The storage requirements of the $O(kn \log n)$ and $O(n \log n)$ algorithms differ. Although both algorithms require storage for a tuples array of n k-dimensional tuples, the $O(kn \log n)$ algorithm requires storage for an n-element temporary array plus k n-element index arrays, whereas the $O(n \log n)$ algorithm requires storage for an $n/2$-element temporary array plus only one n-element index array.

The $O(n)$ median-finding algorithm is somewhat complicated and requires careful implementation to achieve optimum performance. For example, the median-finding algorithm utilizes a sorting algorithm to sort large numbers of five-element arrays. For the initial implementation of the $O(n)$ median-finding algorithm, these small arrays were sorted via the merge sort algorithm that is used for the initial sort of the index array [14]. That merge sort algorithm is not optimized for sorting small arrays and hence resulted in poor performance of the $O(n)$ median-finding algorithm and consequent poor performance of the $O(n \log n)$ k-d tree-building algorithm. Replacing the merge sort algorithm with an insertion sort algorithm [5] that is better suited to sorting small arrays allowed a 30 percent improvement in the performance of the $O(n \log n)$ k-d tree-building algorithm.

5. CONCLUSION

The k-d tree-building algorithm that is proposed in this article achieves a worst-case computational complexity of $O(kn \log n)$ for n points and k dimensions. For $k = 4$, the performance of this algorithm equals the performance of a $O(n \log n)$ k-d tree-building algorithm that employs a $O(n)$ median-finding algorithm. For either algorithm, an improvement in performance by a factor of three relative to single-threaded performance is achieved via parallel execution by eight threads of an Intel quad-core i7 processor that supports concurrent execution of two threads per core.

Source Code

I include C, C++ and Java implementations of the $O(kn \log n)$ and $O(n \log n)$ k-d tree-building algorithms that use OpenMP and Java threads for parallel execution. I also include the R-language code for analysis. The source code for these implementations includes the BSD 3-Clause License.

Acknowledgements

I thank Paul McJones, Gene McDaniel, Joseph Wearing and John Robinson for helpful comments.

6. REFERENCES

[1] G. Adelson-Velskii and E. Landis. An algorithm for the organization of information. *Proceedings of the USSR Academy of Sciences*, 146:263–266, 1962.

[2] G. M. Amdahl. Validity of the single-processor approach to achieving large scale computing capabilities. In *American Federation of Information Processing Societies (AFIPS) Conference Proceedings*, pages 483–485, Reston, VA, 1967. AFIPS Press.

[3] R. Bayer. Symmetric binary b-trees: Data structure and maintenance algorithms. *Acta Informatica*, 1:290–306, 1972.

[4] J. Bentley. Multidimensional binary search trees used for associative searching. *Communications of the ACM*, 18:509–517, 1975.

[5] J. Bentley. Insertion sorts. In *Programming Pearls*, pages 115–116. Addison-Wesley, Reading, MA, second edition, 1999.

[6] M. Blum, R. Floyd, V. Pratt, R. Rivest, and R. Tarjan. Time bounds for selection. *Journal of Computer and System Sciences*, 7:448–461, 1973.

[7] T. Cormen, C. Leiserson, R. Rivest, and C. Stein. Selection in worst-case linear time. In *Introduction to Algorithms*, pages 220–222. MIT Press, Cambridge, MA, third edition, 2009.

[8] J. Friedman, J. Bentley, and R. Finkel. An algorithm for finding best matches in logarithmic expected time. *ACM Transactions on Mathematical Software*, 3:209–226, 1977.

[9] H. Goldstine and J. von Neumann. Coding of some combinatorial (sorting) problems. In A. Taub, editor, *John von Neumann Collected Works: Design of Computers, Theory of Automata and Numerical Analysis*, volume 5, pages 196–214. Pergamon Press Ltd. and the Macmillan Company, New York, NY, 1963.

[10] L. Guibas and R. Sedgewick. A dichromatic framework for balanced trees. In *Proceedings of the*

19th Annual Symposium on Foundations of Computer Science., pages 8–21. IEEE Press, 1978.

[11] N. J. Gunther. Scalability - a quantitative approach. In *Guerrilla Capacity Planning: A Tactical Approach to Planning for Highly Scalable Applications and Services*, pages 41–70. Springer-Verlag, Berlin, first edition, 2007.

[12] V. Havran and J. Bittner. On improving kd-trees for ray shooting. In *Proceedings of the Winter School of Computer Graphics (WSCG)*, pages 209–216, Plzen, Czech Republic, 2002. Science Press.

[13] C. Hoare. Quicksort. *The Computer Journal*, 5:10–15, 1962.

[14] R. Sedgewick. Mergesort. In *Algorithms in C++*, pages 165–166. Addison-Wesley, Reading, MA, 1992.

[15] I. Wald and V. Havran. On building fast kd-trees for ray tracing, and on doing that in O(N log N). In *Proceedings of the 2006 Institute of Electrical and Electronics Engineers (IEEE) Symposium on Interactive Ray Tracing*, pages 61–69. IEEE Press, 2006.

[16] J. Williams. Heapsort (algorithm 232). *Communications of the ACM*, 7:347–348, 1964.

[17] N. Wirth. Finding the median. In *Algorithms + Data Structures = Programs*, pages 82–84. Prentice-Hall, Englewood Cliffs, NJ, 1976.